(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 006 702 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
07.06.2000 Bulletin 2000/23

(51) Int Cl.[7]: **H04L 29/14**, H04L 12/56

(21) Application number: **99309720.3**

(22) Date of filing: **03.12.1999**

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **04.12.1998 US 205084**

(71) Applicant: **NORTEL NETWORKS CORPORATION**
Montreal, Quebec H2Y 3Y4 (CA)

(72) Inventors:
- **Srikanth, Ayikudy K**
  Reading, Massachusetts 01867 (US)
- **Mujeeb, Hamayon**
  Chelmsford, Massachusetts 01824 (US)
- **Meehan, Tom**
  New York, 10002 (US)

(74) Representative: **Dearling, Bruce Clive et al**
Hepworth Lawrence Bryer & Bizley,
Merlin House,
Falconry Court,
Bakers Lane
Epping, Essex CM16 5DQ (GB)

(54) Method and apparatus providing for an improved VRRP (Virtual Router Redundancy Protocol)

(57) A router attached to a network is a member of a virtual router for the network in accordance with VRRP. A port or interface on the router other than a port that attaches the router to the network is specified as critical. VRRP is not enabled for the critical port. (Alternatively, VRRP may be enabled for the critical port to provide routing redundancy for a different network and would not affect the virtual router associated with this network). If the critical port on a master virtual router fails or is otherwise unable to forward packets, a state change is nevertheless triggered within the VRRP finite state machine that exists at the port on the master virtual router for which VRRP is enabled. The change in the VRRP finite state machine triggers the router to transition from the master virtual router to the new backup virtual router. The master virtual router can directly notify the backup virtual router to transition to master virtual router by transmitting a VRRP packet commanding the backup virtual router to become the master virtual router.

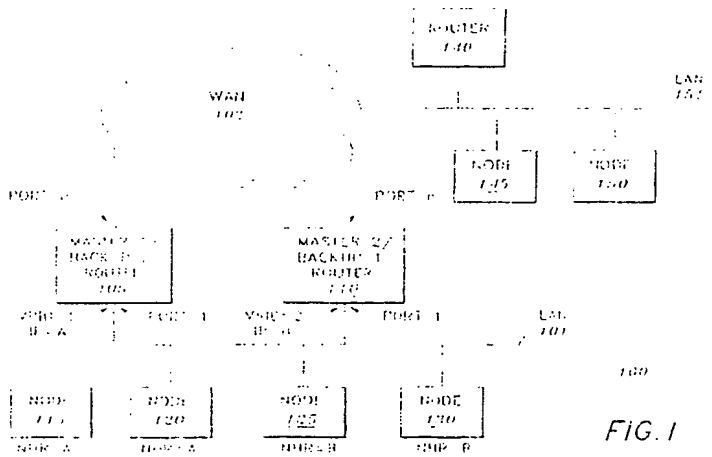

FIG. 1

EP 1 006 702 A2

router for virtual router 1, has an IP address of "A" assigned to port 1, and forwards IP datagrams received at port 1 from the nodes having a statically configured default next hop router IP address of "A".

[0008] The master virtual router periodically transmits advertisements to the backup virtual router(s) on the local network to indicate to the backup(s) that it is still functioning as the master virtual router. If master virtual router 1 fails, the backup virtual router 1 takes over as the new master virtual router 1, providing routing capability for nodes 115 and 120. Since both routers share the same IP address ("IP A") on their respective entry ports, (port 1) and both share the VRRP based media access control (MAC) address on their respective entry ports, no reconfiguration of the static default next hop router IP address is required at each of the nodes that transmit IP datagrams destined for nodes on other IP networks to virtual router 1. Likewise, if master virtual router 2 fails, backup virtual router 2 provides routing for nodes 125 and 130.

[0009] Fig. 2 illustrates a prior art finite state machine 200 for VRRP. An instance of the finite state machine exists for each virtual router in which a VRRP based router is participating. For example, router 105 is a master virtual router in virtual router 1, and a backup virtual router in virtual router 2. Hence, two instances of the finite state machine exist on router 105. In particular, each instance of the finite state machine is associated with a port on VRRP based router. Thus, router 105 has two instances of the VRRP finite state machine associated with port 1 - one instance for virtual router 1, and a second instance for virtual router 2. A VRRP based router begins in initialize state 205, and on a startup event either transition to a master state 215 of a backup state 210, based on its priority. If the router's priority is high, e.g., 255, it transitions to a master state upon the occurrence of a startup event. If the router's priority is less than 255, it transitions to a backup state upon the occurrence of a startup event. In either state, the router returns to the initialize state 205 upon the occurrence of a shutdown event.

[0010] Master virtual routers periodically transmit VRRP advertisements to the appropriate ports of other routers participating in the virtual router, using IP multicast datagrams. If a master virtual router, i.e., a virtual router in master state, receives a VRRP advertisement from a backup virtual router, i.e., a virtual router in backup state, with a priority greater than the master virtual router's priority, or with a priority equal to the master virtual router's priority and a greater IP address (the IP address acts as a tiebreaker), the master virtual router transitions to backup state 210. Conversely, the backup virtual router transitions to master state 215 upon expiration of a master_down_timer, i.e., the backup virtual router fails to receive an advertisement from the master virtual router for a period of time equal to master_down_timer. What is needed is the ability to transition a port associated with the VRRP finite state machine from a master state to a backup state in the event of failure of another port not associated with the VRRP finite state machine.

[0011] The present invention relates to an improvement, or extension of, the virtual routing redundancy protocol (VRRP), as set forth in the Internet Society's Request For Comments 2338 (RFC 2338). In a network that has multiple redundant paths over which packets may be forwarded by VRRP based routers to another network, nodes attached to the network dynamically select which of the routers will forward packets to nodes attached to the other network. VRRP is enabled at the port on each router that attaches the router to the network. A method is described for transitioning responsibility among the routers for routing data packets from the network to another network. One router is initialized to function as a mater virtual router for the network. At least a second router is initialized to function as a backup virtual router for the network. If a port fails on the master virtual router other than the port that attaches the router to the network, even though VRRP is not enabled on the failed port, the master virtual router nevertheless transitions to function as the backup virtual router for the network. The backup virtual router, meanwhile, either times out waiting to receive an advertisement from the master virtual router that would indicate the master virtual router is still functioning as the master virtual router for the network, or the master virtual router, after transitioning to become the new backup virtual router, sends a VRRP packet to the backup virtual router indicating the backup virtual router should become the new master virtual router for the network.

[0012] The present invention is illustrated by way of example and not limitation in the following figures, in which:

[0013] Figure 1 is a diagram of a data communications internetwork.

[0014] Figure 2 is a diagram of a finite state machine for a prior art virtual router redundancy protocol.

[0015] Figure 3 is a diagram of a finite state machine for an improved virtual router redundancy protocol as may be embodied by the present invention.

[0016] Described is an improved virtual router redundancy protocol. In the following description, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one of ordinary skill in the art that the present invention may be practised without these specific details. In other instances, well-known architectures, steps, and techniques have not been shown to avoid unnecessarily obscuring the present invention. For example, specific details are not provided as to whether the method is implemented in a switch as a software routine, hardware circuit, firmware, or a combination thereof. Indeed, with respect to an existing network, operation functionality provided by a software upgrade could be supported by the download of code through the network itself.

the present master virtual router to transition to the back-up virtual router, as depicted at 315.

[0022] In accordance with the invention, there is provided a method for transitioning responsibility for routing data traffic among routers in a network in which a first router is coupled to the network via a first port and a second router is coupled to the network via a first port, the method comprising:-

a) configuring the first router to function as a master virtual router for the network; and
b) configuring the second router to function as a backup virtual router for the network; and
c) if a port on the first router other than the first port fails, transitioning the first router to function as the backup virtual router for the network and transitioning the second router to function as the master virtual router for the network.

Claims

1. A method for transitioning responsibility for routing data traffic among routers in a network in which a first router is coupled to the network via a first port and a second router is coupled to the network via a first port, the method comprising:-

   a) configuring the first router to function as a master virtual router for the network: and
   b) configuring the second router to function as a backup virtual router for the network: and
   c) if a port on the first router other than the first port fails, transitioning the first router to function as the backup virtual router for the network and transitioning the second router to function as the master virtual router for the network, there being provided transitioning means operable to effect the aforesaid transitioning of the first and second routers.

2. A method as claimed in Claim 1 wherein, responsive to the port on the first router other than the first port suffering failure and subsequently becoming re-available, the first router is transitioned to function as the master virtual router for the network and the second router is transitioned to function as the backup virtual router for the network.

3. The method for transitioning responsibility for routing data traffic among routers in a network according to claim wherein 1 or 2,

   transitioning the first router to the backup virtual router and
   transitioning the second router to the master virtual router is responsive to failure of a critical port of the first router.

4. The method of claim 3, wherein responsive to the critical port on the first router suffering failure and subsequently becoming re-available again, the first router is transitioned to function as the master virtual router and the second router is transitioned to function as the back-up virtual router.

5. The method of any preceding claim, wherein configuring the first router as a master virtual router comprises configuring a priority for the first router, the priority determining that the first router becomes the master virtual router, and network preferably operating a virtual routing redundancy protocol (VRRP) and wherein the priority is a VRRP priority.

6. The method of any preceding claim, wherein configuring the second router as a backup virtual router comprises configuring a priority for the second router, the priority determining that the second router becomes the backup virtual router, the network preferably operating a virtual routing redundancy protocol (VRRP) and wherein the priority for the second router is a VRRP priority.

7. The method for transitioning responsibility for routing data traffic among routers in a network according to any preceding claim, further comprising transmitting a message, upon failure of the port on the first router, from the first router to the second router that instructs the second router to become the new master virtual router, and transitioning the first router to become the new backup virtual router.

8. The method of claim 7, wherein transmitting a message from the first router to the second router comprises transmitting a virtual routing redundancy protocol message from the first router to the second router.

9. The method of any one of claims 1 to 6, further comprising: at the second router, periodically receiving a multicast advertisement generated from the first router else causing the second router to become the new master router after a predetermined time has elapsed in which the multicast advertisement has not been received by the second router.

10. A computer program element comprising computer program code means to make a computer-controlled router execute procedure to perform the method steps of any preceding claim.

11. The computer program element of claim 10, embodied on a computer readable medium.

12. Electronic signals representing instructions or statements to make a computer-controlled router execute procedure to perform the method steps of

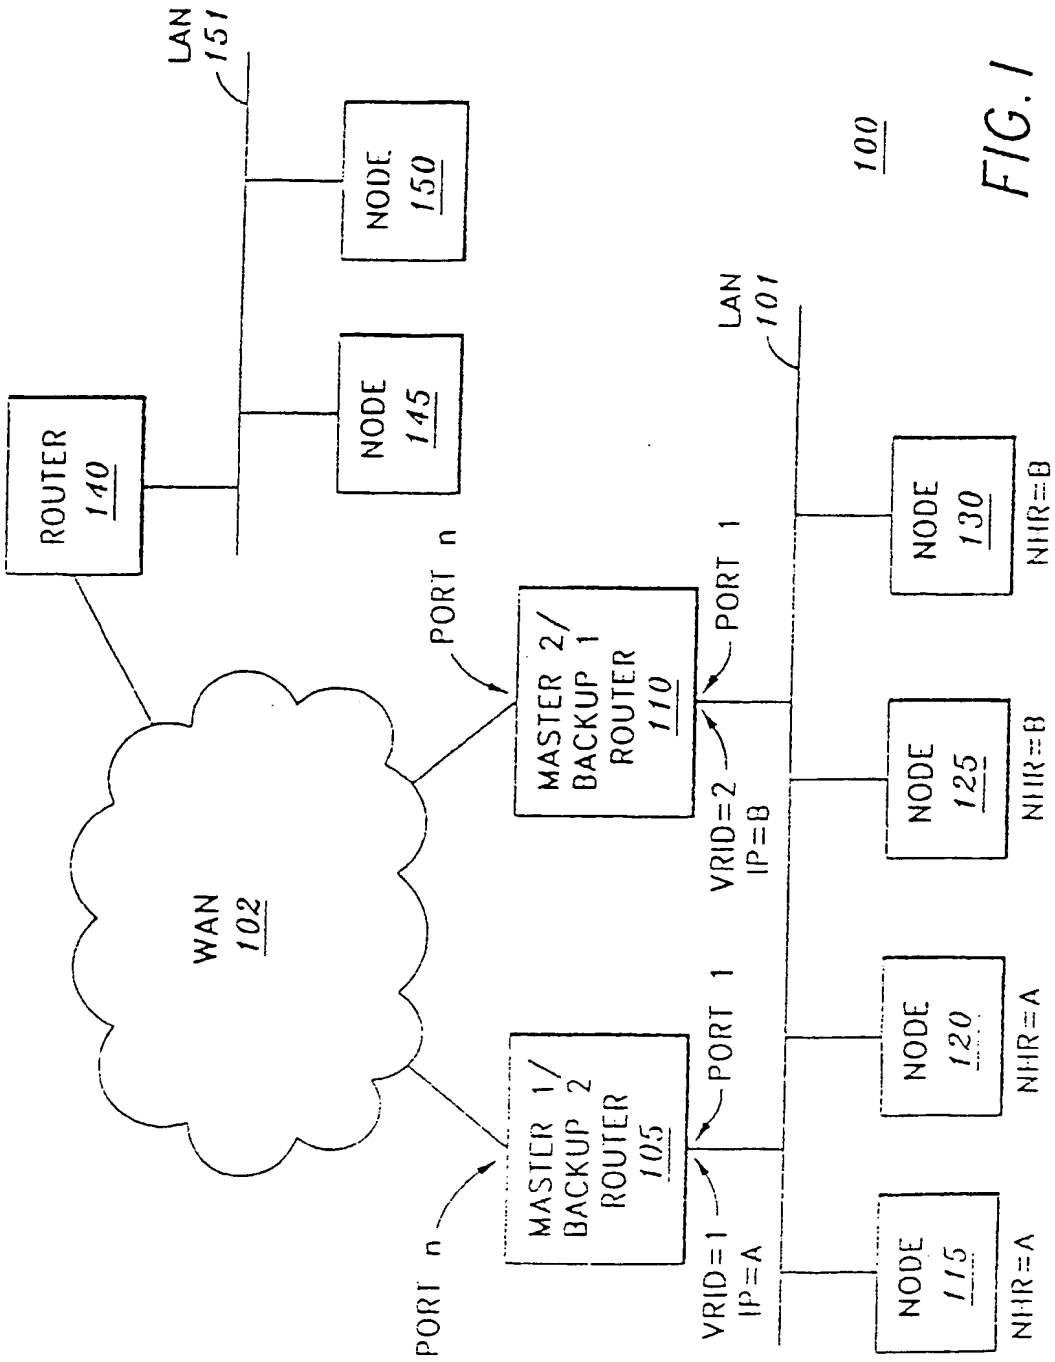

FIG. 1

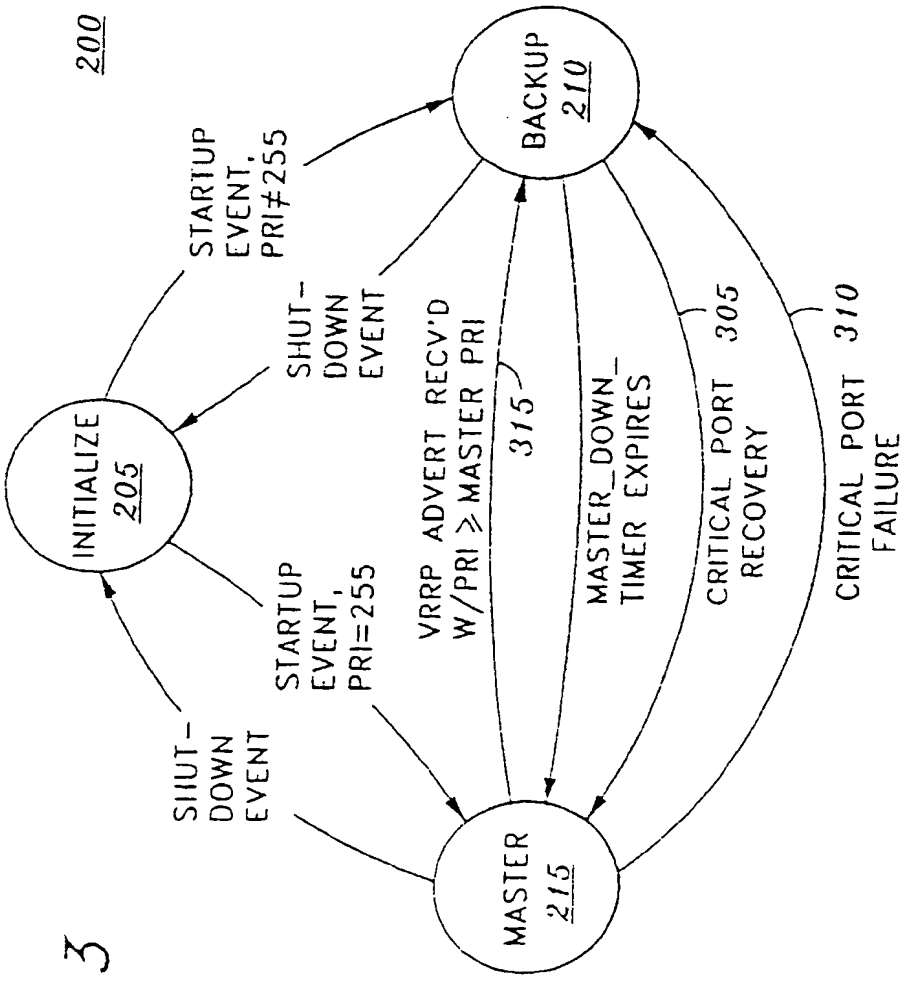

FIG. 3

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 006 702 A3

(12) EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
14.06.2000 Bulletin 2000/24

(43) Date of publication A2:
07.06.2000 Bulletin 2000/23

(21) Application number: **99309720.3**

(22) Date of filing: **03.12.1999**

(51) Int Cl.[7]: **H04L 29/14**, H04L 12/56

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **04.12.1998 US 205084**

(71) Applicant: **NORTEL NETWORKS CORPORATION**
Montreal, Quebec H2Y 3Y4 (CA)

(72) Inventors:
- **Srikanth, Ayikudy K**
  Reading, Massachusetts 01867 (US)
- **Mujeeb, Hamayon**
  Chelmsford, Massachusetts 01824 (US)
- **Meehan, Tom**
  New York, 10002 (US)

(74) Representative: **Dearling, Bruce Clive et al**
Hepworth Lawrence Bryer & Bizley,
Merlin House,
Falconry Court,
Bakers Lane
Epping, Essex CM16 5DQ (GB)

(54) Method and apparatus providing for an improved VRRP (Virtual Router Redundancy Protocol)

(57) A router attached to a network is a member of a virtual router for the network in accordance with VRRP. A port or interface on the router other than a port that attaches the router to the network is specified as critical. VRRP is not enabled for the critical port. (Alternatively, VRRP may be enabled for the critical port to provide routing redundancy for a different network and would not affect the virtual router associated with this network). If the critical port on a master virtual router fails or is otherwise unable to forward packets, a state change is nevertheless triggered within the VRRP finite state machine that exists at the port on the master virtual router for which VRRP is enabled. The change in the VRRP finite state machine triggers the router to transition from the master virtual router to the new backup virtual router. The master virtual router can directly notify the backup virtual router to transition to master virtual router by transmitting a VRRP packet commanding the backup virtual router to become the master virtual router.

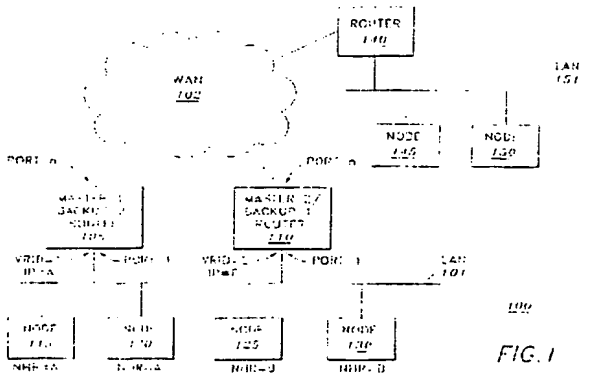

FIG. 1

EP 1 006 702 A3

ANNEX TO THE EUROPEAN SEARCH REPORT ON EUROPEAN PATENT APPLICATION NO. EP 99 30 9720

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

10-04-2000

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5815668 A	29-09-1998	JP 2708009 B JP 8256173 A	04-02-1998 01-10-1996
US 5473599 A	05-12-1995	AU 2359995 A WO 9529544 A	16-11-1995 02-11-1995
US 5835696 A	10-11-1998	NONE	

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office. No. 12/82